

UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP010531

TITLE: Lessons Learned in Developing and
Validating Models of Visual Search and Target
Acquisition

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Search and Target Acquisition

To order the complete compilation report, use: ADA388367

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, ect. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP010531 thru ADP010556

UNCLASSIFIED

Lessons Learned in Developing and Validating Models of Visual Search and Target Acquisition

¹Theodore J. Doll and ²Richard Home

¹Electro-Optics, Environment, and Materials Laboratory
Georgia Tech Research Institute
Georgia Institute of Technology
Atlanta, GA 30332-0841, USA
Phone: (+) 404 894 0022
Fax: (+) 404 894 6199 / 6285
E-mail: ted.doll@gtri.gatech.edu

²Defence and Evaluation Research Agency
Malvern, Worcester, WR14 3PS, UK
Phone: (+) 44 1684 896 933
Fax: (+) 44 1684 896 714
E-mail: rhome@dera.gov.uk

1. SUMMARY

Some shortcomings of past and current approaches for modeling human visual search and target acquisition (STA) are discussed. The effects of complex pattern perception, visual attention, learning, and cognition on STA performance are particularly emphasized. The importance of these processes is explained and approaches are suggested for modeling them. Guidelines are also provided for testing and validating models of visual search and target acquisition. These guidelines take into account the roles of pattern perception, visual attention, learning, and cognition in STA performance. The present paper also presents and compares alternative approaches to field testing for the purpose of model validation.

Keywords: search, target acquisition, perception, attention, learning, cognition, validation

2. INTRODUCTION

The military spends millions of dollars annually to build large-scale, system-level simulations of weapons and related systems. These simulations enable their users to understand how the systems will perform under conditions that would be impossible or extremely costly to produce in the real world. However, very little money is spent on system-level simulation of the one system that is key to all military operations – the human visual system.

System-level simulations of human vision could be useful in setting performance standards for both the naked eye and all types of sensors and systems in which the final judgement or interpretation is made by a human observer. System-level simulations of human vision would also lead to more accurate design requirements for sensors and camouflage, concealment, and deception (CCD) systems. A better understanding of the human visual system would also provide insights into how best to test and validate models of search and target acquisition (STA) performance.

Until recently, attempts to build general models of human observer target acquisition performance have met with only limited success. By the term “general”, we mean models that accurately predict the detectability of (at least) military targets as viewed through a wide variety of sensors in a wide variety of backgrounds, without the need for calibration in each new situation. The difficulty no doubt stems in part from the inherent complexity of human perception and performance – but also in part from the manner in which the problem has been approached by the military R&D community.

Military-sponsored STA modeling has traditionally followed either of two approaches: (1) physics-based, or (2) simple models of human visual performance that emphasize only a part of the neural “machinery” involved in human STA. The physics-based approach is based on the idea that simply matching the target signature to the background clutter will suffice to deny detection. In spite of decades of research, this approach has failed. The reason is that no one has been able to determine to which aspects of the background clutter it is necessary to match to the target. It has proven impossible to match targets to all aspects of background clutter because clutter characteristics change over and within scenes (i.e., clutter is non-stationary).

Modeling efforts following the second approach – modeling only a limited part of the visual system – have typically emphasized the basic sensitivity of the eye to light, or at best, the basic spatio-temporal contrast sensitivity of the visual system. They typically pay scant attention to the important roles of complex pattern perception, visual attention, learning, and cognition in STA performance. Thus, they model only a limited part of the visual system. This state of affairs has occurred, in large part, because there has not been wide-spread understanding of the attentive, perceptual, and cognitive aspects of visual

performance and the role of learning in the military R&D community.

There is, however, a widening awakening to the role of attention, perception, cognition, and learning. Some of the papers in this conference attest to that fact. In his abstract, Al Ahumada remarks that "learning and memory components are required for a model that can accurately predict human detection in unpredictable backgrounds." In discussing shifts of attention during search, John Findlay suggests that "eye movements are programmed on the basis of a spatial salience map with both excitatory and inhibitory influences reaching it from feature maps", and "flexibility in search is provided through learning mechanisms." Commenting on the role of perceptual organization in contour and texture segregation, Wilson Geisler notes that "evidence suggests that more sophisticated models incorporating perceptual organization mechanisms will be required to predict human texture and contour segregation performance."




Figure 1. Tank and background with identical first-order statistics.

The cost of ignoring attention, perception, cognition, and learning is that the models developed have limited scope, and must be empirically calibrated for each new sensor technology, background environment, CCD technique, and level of observer experience. In the remainder of this paper we will explain which aspects of attention, perception, cognition, and learning we believe are most important and why they must be modeled in order to predict STA performance accurately and generally. We will also describe the manner in which these processes have been implemented in one model of human search and detection performance – the Georgia Tech Vision (GTV) model.¹

Unfortunately, including all the relevant visual processes leads to very complex models that are difficult to validate, as Richard Hecker notes in his abstract. However, we disagree with Dr. Hecker's implication that higher perceptual processes like recognition, identification, and search can be eliminated from a model and still have it generate accurate predictions. We will therefore also discuss requirements for model testing and validation that take into account these higher level processes.

3. ROLES OF ATTENTION, PERCEPTION, COGNITION, AND LEARNING IN STA PERFORMANCE

3.1. Perceptual Organization

Walker and McManamey² point out that first-order statistics do not provide information about the spatial structure of an image. First-order metrics include the mean and standard deviation, as well as some less well-known metrics like the Doyle metric and measures of histogram similarity. The tank and the background shown in Figure 1 have identical means and standard deviations, and they're also identical in terms of the Doyle metric and histogram similarity. But they differ in terms of the arrangement, or spatial structure, of the pixels of various gray-scale values. The fact that the tank is clearly detectable from the background demonstrates that first-order statistics are not sufficient.

To account for the detectability of this target we must consider second-order metrics. The gray level co-occurrence matrix (GLCM) is one second-order metric; others include the correlation length and the co-occurrence matrix³. Both of these quantify the correlation between gray-scale values various numbers of pixel locations apart. Although the GLCM, correlation length, and co-occurrence matrix capture some of the properties that contribute to detection, they don't capture all of them. There are texture differences that humans can distinguish, but to which GLCM and correlation length metrics are insensitive.

The image in Figure 2a contains a texture irregularity that human observers can detect (note center bar-shaped region in the center of the image). However, most metrics and models of vision cannot detect this irregularity⁴. This is true of both single-stage, oriented linear-filter models and metrics like the correlation length, co-occurrence matrix, and GLCM. The reason for this is that the entire pattern is made up of the same texture elements – lines of the same length at different orientations. In addition, the probabilities of gray-level transitions from point to point are virtually the same in the center "irregularity" and the surround regions of the image. What distinguishes the center region is not the texture elements themselves, but their relationship to one-another. Note around the center region, that there are abrupt transitions in the relative orientations of the line elements. In the background, by contrast, the orientations of adjacent line elements change only slowly.




Figure 2a. Input image with texture transition near center.




Figure 2b. Output of a single-stage, simple cortical cell, filter model.




Figure 2c. Output of two-stage, complex cortical cell, vision model.

Another way of thinking about this pattern is that the center region is defined by a texture transition. In order to detect these subtle texture transitions, a vision model must have a second filtering stage, which models the outputs of complex cortical cells. Figure 2b shows the output produced by a model with only simple cortical cell (single-stage) filters, for the input in Figure 2a. Note that there is no differential signal that distinguishes the center irregularity.

The GTV model has a second filtering stage, as shown in Figure 3. Each first-stage output is routed to multiple second-stage, spatial-frequency band-pass filters. Depending on the version of GTV run, the second-stage filters may also be orientation-selective. The second stage filters smooth the outputs of each first-stage over regions of

various sizes and orientations. This smoothing serves to identify the extent, or boundaries, of each type of texture identified by the first-stage filters. By comparing these boundaries, GTV can identify texture boundaries, as shown in Figure 2c.

But is the detection of such subtle texture transitions relevant to real-world CCD problems? Figure 4a shows a texture transition that might occur with a perfectly camouflaged vehicle positioned against a background of vegetation. When the vehicle is repositioned, there will be a phase mismatch between the texture of the vegetation and the camouflage pattern on the vehicle. Figure 4b shows a GTV output for this pattern, after the model is trained to detect similar phase-mismatched targets.

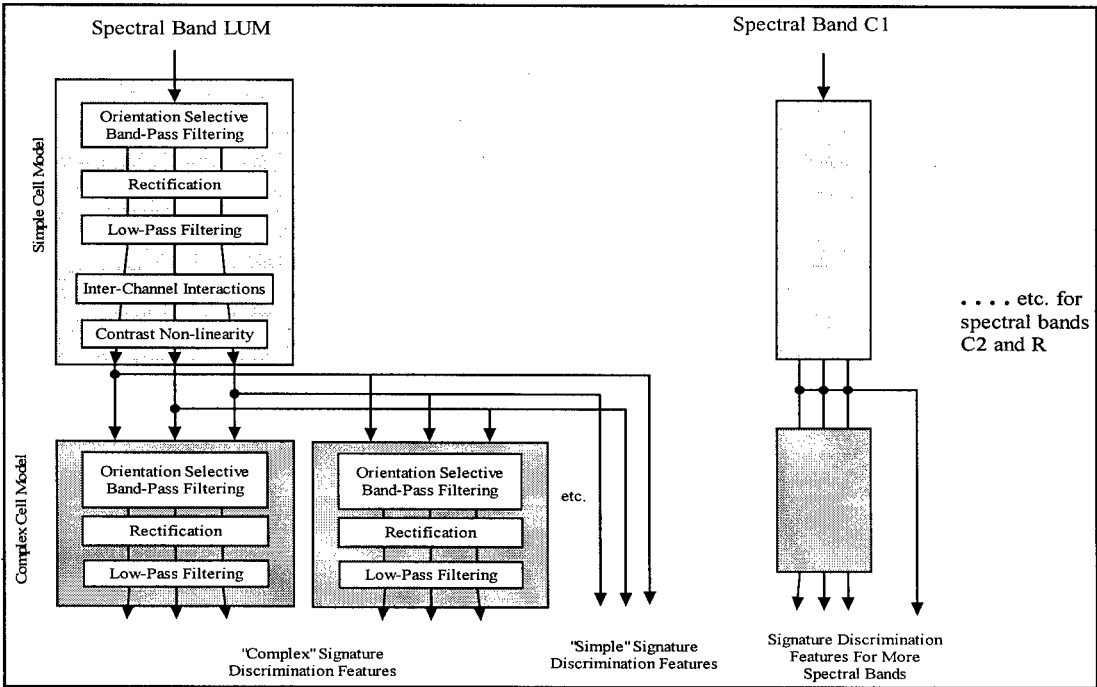


Figure 3. Schematic of GTV two-stage filter process, simulating complex cortical cell outputs.




Figure 4a. Object perfectly matched in pattern and chromaticity to background.

3.2. Attention and Search

There is substantial evidence that eye movements (saccades) during visual search are guided by preattentive (unconscious) processing of pattern information in peripheral vision. For example, recordings of eye movements over structured scenes reveal that the eye fixates on features such as edges and corners that are more likely to convey information than are plain surfaces.⁵ In reading, the eyes of proficient readers search out larger words, which convey a higher degree of meaning than do small words, such as articles⁶. Visual search proficiency has even been used as a measure of peripheral visual acuity.⁷⁻⁹

The implications of this are:

- That clutter (i.e., the input scene) drives visual search.
- That successful search is a prerequisite for detection.
- The eyes fall on those objects that are most conspicuous.
- The assumption, often made in vision models, that search is random is false.

The first line of self-protection is not to be noticed in the first place, that is, to deny visual search. It's generally easier to prevent an observer from locating a target than it is to deny detection once he's looking directly at the target. This is especially true in medium- to high-clutter environments.

Explicit modeling of the effect of clutter on visual search is therefore necessary to accurately predict target acquisition. High clutter in an image reduces probability of locating the target, given limited search time. The GTV model predicts the fixation locations based on the spatial and temporal contrast of objects in the input image. A salience map is




Figure 4b. Output of GTV identifying pattern in Figure 4a.

generated by using multiple-channel, quasi-linear filtering mechanisms. This map also serves as a basis for segregating the input scene into areas of interest for further (attentive) processing.

Another aspect of search that affects target acquisition is the temporal sequence of eye fixations in a scene. A wealth of data shows that human observers tend not to immediately re-fixate on objects when inspecting a scene.¹⁰⁻¹² The GTV model includes a systematic search routine which simulates the fact that observers tend to disregard objects that they have recently fixated and determined not to be targets. Thus, if an object has a high probability of fixation on one glimpse, and it is determined not to be a target, it will be less likely to be fixated on the next glimpse. The systematic search routine also simulates the tendency of observers to eventually re-fixate objects that were previously fixated and found not to be targets. Fixation probabilities that were initially high and decreased tend to recover (increase again) after a number of glimpses. The recovery time depends on the number of blobs in the field of view. This is consistent with empirical studies of visual search.

3.3. Selective Attention and Perceptual Learning

There are at least two aspects of attention that are important to STA performance. One – the mechanism that determines eye fixations and preattentive shifts in visual attention – was discussed in the last section. A second concerns the nature of the visual features that contribute to preattentive “pop-out” of objects and whether those features are subject to modification through learning. In the 1970's, Ann Treisman and her colleagues argued that preattentive processing and selection occur only for objects that are uniquely distinguished by a single perceptual dimension, such as size, color, shape, and luminance. However, Jeremy Wolfe and his colleagues later showed that given sufficient practice, observers could preattentively identify

objects based on conjunctions of perceptual dimensions (e.g., find the red circle in a background of blue circles, blue squares, and red squares). Neisser and others have shown that, given enough experience with the stimulus, observers can reach a point where complex combinations of features support pop-out. For example, Neisser found that after extensive training, observers can learn to rapidly pick out a target letter that is very similar to the background clutter, e.g., a "K" in a field of Es, Hs, Ts, Ls, and Fs. Schneider and his colleagues have studied the development of "automaticity" or preattentive processing in letter search tasks. They showed that letters that are consistently "mapped" as one of a set of targets (as opposed to sometimes being targets and sometimes distractors) eventually become automatically processed after extensive practice.

weighting routine is highly effective in rejecting clutter, as shown in Figure 5, and it allows the model to simulate the performance of experienced human observers.

3.4. Other Cognitive Processes that Affect STA

Another aspect of cognition that affects search and target acquisition is perceptual decision making. Target acquisition is not simple signal-to-noise ratio threshold process, but involves decision-making. Signal detection theory describes observers' ability to trade-off detections versus false alarms. These trade-offs can distort the relative probabilities of detection in task of differing difficulty¹³. For example, we have previously reported that human observers tend to shift their decision criterion as the difficulty of the detection task changes. For example, in

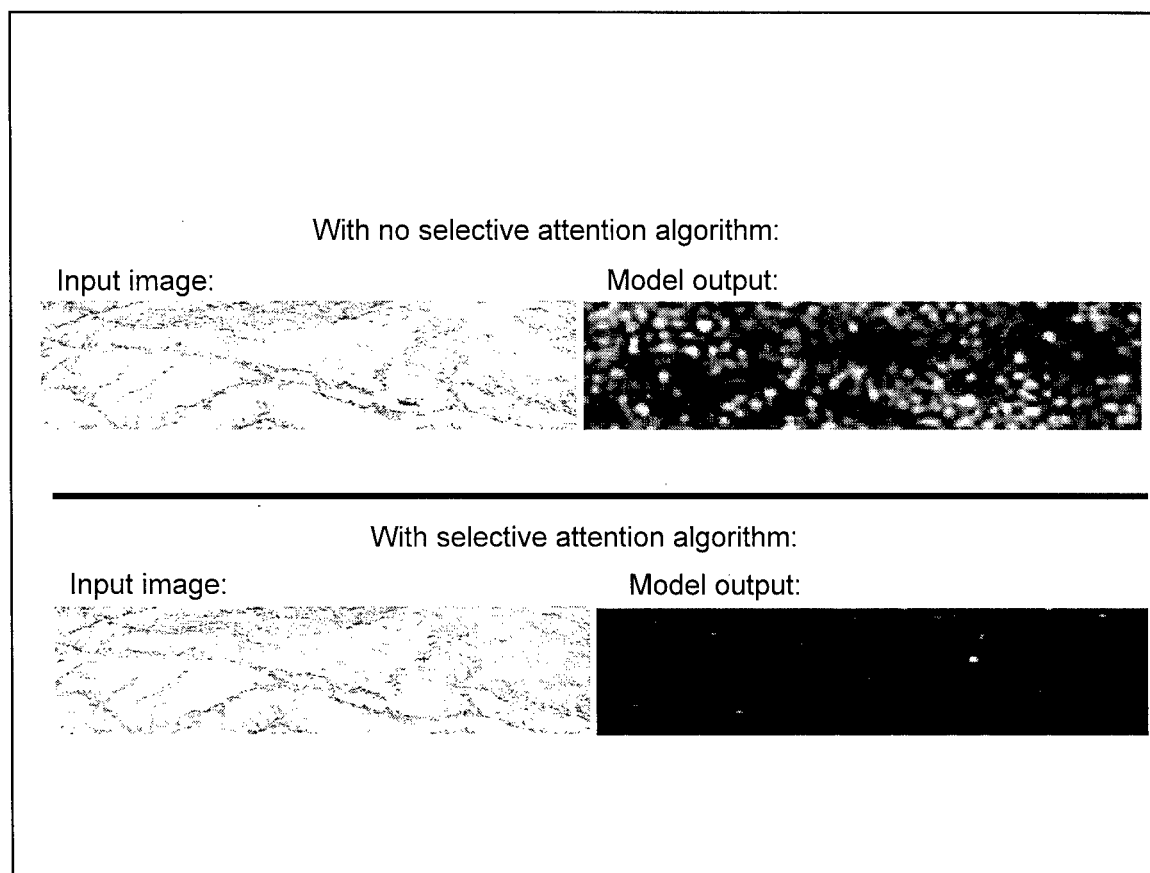


Figure 5. Clutter rejection performance of the GTV model.

After extensive practice, military observers are often able to immediately pick out targets in cluttered scenes that novice observers must search for painstakingly. They have evidently learned to preattentively process the target. It is therefore important to model the effect of learning on pop-out and visual search performance. One way of doing this is to differentially weight the filter-channel outputs before pooling them into a single salience map. The weights would be designed to amplify channel outputs typical of the target, and attenuate channel outputs typical of background clutter. The GTV model uses this method, employing a discriminant analysis routine to compute the weights. The

low clutter conditions the observer may adopt a relatively high decision likelihood ratio criterion, β . But when faced with high clutter, the same observers tend to relax β . This has the effect of allowing them to increase their probability of detection at the cost of a higher false alarm probability, as illustrated in Figure 6. This perceptual decision tradeoff process can have considerable impact on measured probabilities of detection.




Figure 6. ROC curve showing shift in observer criterion with task difficulty.

The GTV model uses signal detection theory in two ways. When there are multiple “blobs” or areas of interest in the field of view, a decision must be made as to which blob the eyes will saccade to next. The extreme detector model is used to make this decision. The choice of blobs for the next saccade is highly non-linear – even though one blob may have just slightly greater spatio-temporal contrast energy than the others, its probability of fixation will be much larger. The metric used to describe each blob is actually a power function of its spatio-temporal contrast energy.

The GTV model also uses signal detection theory to decide whether or not the blob currently being fixated is a target. The spatio-temporal contrast metric for the current blob is compared to the distributions of the same metric for targets and clutter objects encountered during training. It is predicted that the observer says “yes, the blob is a target” when the ratio of the probability densities of the target to clutter distributions exceeds the criterion value of the decision likelihood ratio, β .

4. REQUIREMENTS FOR TESTING AND VALIDATING STA MODELS

There are a number of requirements for the design and conduct of successful validation tests that derive from an understanding of human vision and visual cognition. Although many investigators will be familiar with these requirements, one or more of the requirements have not been met in almost all STA model validation efforts. Exposition and discussion of requirements can therefore benefit the STA community.

- Since the sensitivity of the human visual system depends on the luminance level and chromaticity of the input scene, input images must be photometrically and colorimetrically calibrated. Some issues of color calibration are discussed by Rogers and Thomas¹⁴.
- Since the human visual system has high acuity only in small portion of the visual field (i.e., the fovea), the likelihood that a target is foveated is an important determinant of overall detection

probability. The larger the observer’s field of view, the less likely it is that any given target will be foveated (assuming constant magnification). It is therefore important that the apparent field of view (AFOV) of the imagery used to test models be the same as the AFOV that observers used in the experiment whose results are to be matched.

- Simply instructing observers to make their responses indicative of a given level of processing (e.g., locate “areas of interest” without full detection or recognition) does not guarantee that they limit their processing to that level. If the observers are given enough time, they generally perform higher levels of processing (e.g., recognition or identification) before reporting the location of an area of interest. Even if exposure time is limited, observers may perform additional processing on the persisting iconic memory of the target. Observer validation experiments should therefore use brief image exposures followed by a noise mask pattern in order to limit processing.
- If the model under test requires training, the target and background images given the model during training must adequately sample the same target and background features that will be present in the test imagery.
- Two possible scenarios must be considered in determining the spatial resolution of imagery used to test a model: (a) the resolution in the observer test was limited by a display and/or sensor, or (b) the resolution in the observer test was limited only by the human eye, e.g., observers viewing targets with the naked eye or DVO in clear conditions. In the first case, the images submitted to the model must be filtered to simulate the MTF of sensor/display system. In the second case, the images provided as inputs to the model must have resolution at least as great as that of the human visual system. They must therefore be captured by a sensor whose resolution exceeds that of the human visual system.
- The temporal up-date rate of the imagery should be at least the Nyquist frequency of the highest rate of temporal modulation in the scene. Alternatively, frame rate can be set to the highest temporal cut-off frequency of the human eye. This last quantity will depend on the intensities, spatial frequencies, and chromaticities in the scene and the viewing conditions.

It should be noted that these requirements are a product of the complexity of human observers’ visual performance – not a consequence of the complexity of any model. They therefore apply regardless of whether one is testing a simple or a complex model.

5. ALTERNATIVE APPROACHES FOR FIELD TESTING AND MODEL VALIDATION

The process of validating search and detection models or metrics is expensive and time-consuming. It is therefore worth considering some of the alternative approaches available and the advantages and potential pitfalls of each. We contrast three different approaches here, all of which

in backgrounds collected from the field. This is Approach B in Table 1, and the approach used by TNO for the DISTAFF data set. This approach does not eliminate the camera dynamic range problem, but ensures that both the observers and the STA model are subject to the same effects in this regard. However, this approach still suffers from other disadvantages (which are also present in

Table 1. Alternative approaches for field testing and STA model validation.

Approach A Observer test in the field	<ul style="list-style-type: none">• Collection of imagery of targets in backgrounds, ground truth, ambient illumination, meteorological data, and calibration data in field with high resolution camera• Observer test in field viewing targets through DVO device• Field imagery and calibration data submitted to STA model to generate predictions• Model predictions compared to observer performance in field
Approach B Observer test in the laboratory with field imagery	<ul style="list-style-type: none">• Collection of imagery of targets in backgrounds, ground truth, ambient illumination, meteorological data, and calibration data in field with high resolution camera• Observer test in the laboratory by displaying imagery from field test• Field imagery and calibration data submitted to STA model to generate predictions• Model predictions compared to observer performance in laboratory
Approach C Use and validation of synthetic imagery Observer test in the laboratory with synthetic imagery	<ul style="list-style-type: none">• Collection of imagery of background only, ground truth, ambient illumination, meteorological data, and calibration data in field with high resolution camera• Measurement of Bi-directional Reflectivity Distribution Function (BRDF) of target paints• Synthetic target generated and inserted in calibrated background imagery from field• Synthetic imagery validated by comparing it to field imagery• Observer test in the laboratory with validated synthetic imagery• Synthetic imagery submitted to STA model to generate predictions• Model predictions compared to observer performance in laboratory

involve collection of imagery from the field and psychophysical tests with human observers in either the field or a laboratory. The three approaches are summarized in Table 1.

The conventional and most obvious approach is to collect both observer data and imagery to submit to the STA model in the field. This is Approach A in Table 1. One of the major disadvantages of Approach A is that it is difficult to control observer tests in the field. The field of view, exposure time, time of day, and cloud shadows experienced by observers all must be the same as those in the imagery collected for submission to the STA model. Moreover, the observers must be shielded from acoustic and social cues that would affect their STA performance. Another serious problem with Approach A is that no camera can reproduce the full range of colors and intensities that the observers experience in the field. Very high signals (e.g., from specular reflections) will exceed the dynamic range of the camera (i.e., saturate). If the camera gain is set lower, then low signals (e.g., in shadowed areas of the scene) will fall below the sensitivity threshold of the camera and these areas will appear black in the image.

One possible solution to these problems is to do the observer testing in the laboratory using imagery of targets

Approach A). For one, it is expensive and time-consuming to deploy real targets in the field in a controlled manner. The very act of deploying them also produces extraneous detection cues, such as vehicle tracks.

Capturing temporal effects is also a problem in both approaches A and B. If one wants to capture important effects of target motion (relative to the background, or motion of parts of the target relative to the whole), then the problems of field deployment and control are compounded. For example, the rate and pattern of motion of a vehicle over rough terrain may be an important detection/recognition cue. Shadows produced by clouds and the motion of helicopter rotor blades are other temporal effects that can greatly influence detection. Capturing these motion effects in imagery requires a very high frame rate, and results in a huge amount of imagery that must be stored and calibrated.

Extraneous cues from target deployment can be eliminated and temporal effects controlled by using synthetic imagery for both the observer test and as input to the STA model. This approach is used in the VISEO system¹⁵, and is shown as Approach C in the above table. This is a two-step approach – first synthetic imagery is generated and validated, and then the STA model is validated using the

synthetic imagery. The VISEO system generates backgrounds using one or more spectral bands of measured background imagery, depending on the type of sensor being simulated. The spectral bands range from the visible to LWIR. The database is calibrated, and the algorithm for combining bands has been validated.¹⁶ The VISEO system also has a library of approximately 75 high-fidelity ground and air targets, most of which have been validated in the visible and/or IR bands. With the VISEO system, one can generate imagery at any desired frame rate in order to capture high temporal-frequency effects. With VISEO, one need not generate the imagery for the whole set of test conditions at one time. Imagery can be generated for selected conditions, submitted to the STA model to generate predictions, and then archived. Another advantage of the VISEO system is that radiation from the target model is not limited by any camera or sensor system. One can therefore model specular reflections from the target, for example, and evaluate their effect on detectability by submitting the resulting scene data directly to the STA model.

It is clear that Approach A has serious shortcomings – due both the difficulty of controlling observer test in the field and sensor dynamic range limitations. However, approaches B and C both have advantages for certain types of applications. With VISEO, there is no need to deploy and control targets during field imagery collection, and one can more easily evaluate temporal effects and specular reflections. However, one must build high-fidelity models of the targets, if they are not already in the VISEO database.

6. REFERENCES

1. Doll, T. J., McWhorter, S. W., Wasilewski, A. A. and Schmieder, D. E., "Robust, sensor-Independent target detection and recognition based on computational models of human vision", *Optical Engineering*, 37, (7), pp. 2006-2021, 1998.
2. Walker, G. W., and McManamey, J. R., "The importance of second-order statistics for predicting target detectability", In: *Proceedings SPIE*, vol. 1967, pp. 308-319, 1993.
3. Rotman, S. R., Hsu, D. Cohen, A., Shamay, D., and Kowalczyk, M. L., "Textural metrics for clutter affecting human target acquisition", *Infrared Physics & Technology*, 37, pp. 667-674, 1995.
4. Northdurft, H. C., "Texture segmentation and pop-out from orientation contrast", *Vision Research*, 31, pp. 1073-1078, 1991.
5. Gould, J. D., "Looking at pictures", In: R. A. Monty and J. W. Senders (Eds.), *Eye movements and psychological processes*, Hillsdale, N. J.: Lawrence Erlbaum Associates, 1976.
6. Rayner, K., "Foveal and parafoveal cues in reading", In: J. Requin (Ed.), *Attention and performance* (Vol. 7), Hillsdale, NJ: Lawrence Erlbaum and Associates, 1978.
7. Bellamy, L. J. and Courtney, A. J., "Development of a search task for the measurement of peripheral visual acuity", *Ergonomics*, 24, pp. 497-509, 1981.
8. Erikson, R. A., "Relation between visual search time and visual acuity", *Human Factors*, 6, pp. 165-178, 1964.
9. Johnston, D. M., "Search performance as a function of peripheral acuity", *Human Factors*, 7, pp. 528-535, 1965.
10. Arani, T., Karwan, M. H. and Drury, C. G., "A variable-memory model of visual search", *Human Factors*, 26, pp. 631-639, 1984.
11. Humphreys, G. W. and Muller, H. J., "Search via Recursive Rejection (SERR): A Connectionist Model of Visual Search", *Cognitive Psychology*, 25, pp. 43-110, 1993.
12. Nicoll, J. F., *A Search for Understanding: Analysis of Human Performance on Target Acquisition and Search Tasks Using Eyetracker Data*, IDA Paper P-3036, Institute for Defense Analyses, 1995.
13. Doll, T. J. and Schmieder, D. E., "Observer false alarm effects on detection in clutter", *Optical Engineering*, 32, pp. 1675-1684, 1993.
14. Rogers, G. A. and Thomas, D. J., *Improved color image calibration*, Report ???, U.S. Army Combat Systems Test Activity, Aberdeen Proving Ground, MD. U.S. Army Tank-Automotive Research Development and Engineering Center, 1997.
15. Doll, T. J., McWhorter, S. W., Schmieder, D. E., Hetzler, M. C., Stewart, J. S., Wasilewski, A. A., Owens, W. R., Sheffer, A. D., Galloway, G. L., and Harbert, S. D., "VISEO" System for Camouflage /LO Design based on Computational Vision Research, In: *Proceedings of the HAVE FORUM Low Observables Symposium*, Eglin AFB, FL, Report No. WL-TR-97-6003, 1997.
16. Doll, T. J., McWhorter, S. W., Hetzler, M. C., Stewart, J. M., Wasilewski, A. A., Schmieder, D. E., Owens, W. R., Sheffer, A. D., Galloway, G. L., and Harbert, S. L., *Visual / Electro-Optical (VISEO) Detection Analysis System: Final Report*, Report No. USAAMCOM TR 98-D-19, Atlanta, GA: Georgia Tech Research Institute, 1997.